

REMARKS

According to the Final Office Action of June 6, 2002, claims 1-8, 12, 14-15, 21, 23, 25-27, 31, and 34 stand rejected. Claims 9-11, 13, 16-20, 22, 24, 28, 29, 33, 35, and 36 stand objected to.

This is the starting point of this RCE, and the following remarks address themselves to the rejections in the parent application.

According to the Office Action summary page, claim 35 is objected to. In the Detailed Remarks, claim 35 is rejected under 35 USC 103. Thus, there is a discrepancy concerning claim 35. Taking the more disadvantageous position, claim 35 is treated here as having been rejected.

Claims 1, 5-8, 25-27, 31, 32 and 34 were rejected under 35 USC 102 as being anticipated by Afek et al, US Patent 5,748,901 (Afek). Claims 2-4, 21, 23, and 35 were rejected under 35 USC 103 as being unpatentable over Afek in view of Mitra et al, US Patent 6,331,986 (Mitra). Claims 12 and 14-15 were rejected under 35 USC 103 as being unpatentable over Afek in view of Szentesi, US Patent 5,844,886. Applicant respectfully traverses.

Preliminary Observation

In the previous response, applicant has discussed a number of different features of claim 1 that demonstrate claim 1 to be not anticipated by Afek. In the current Office Action, the Examiner rejects claim 1 with text that is word-for-word identical to the text that was previously used to reject claim 1, and in the “*Response to Arguments*” section challenges some – but not all – of the arguments set forth by applicant.

Since some of the arguments were not challenged by the Examiner, it appears that claim 1 should have been held allowable. Moreover, applicant respectfully disagrees with at least some of the challenges, as discussed below. In applicant’s view, the inescapable conclusion is that claim 1 is not anticipated by Afek.

Operation of the Afek system

To the best of applicant’s understanding, the Afek system operates as follows.

A source sends cells at some Allowed Cell Rate (ACR). Interspersed among the cells that carry data there are Resource Management (RM) cells that contain various

fields. One of those fields is the Explicit Rate (ER) field. Another of those fields is the CCR field, which contains the ACR of the session to which the cell belongs. As an RM cell passes through a link and enters a switch, the value of its ER field is replaced, in accord with the FIG. 1 equation, with the lower of:

- (a) a utilization_factor times MACR of the link,
- (b) the current value of the ER field, or
- (c) twice the current cell rate (CCR).

The MACR is a fair share parameter (col. 2, line 48) with a value that is computed for each link that is traversed, based on the unused capacity, Δ , on the traversed link. The computation of MACR of the link is disclosed in detail in Afek's FIG. 1.

It is clear, however, that since the link's MACR is related to Δ , it is related to all other sessions that share the link.

It is also clear that there is a MACR computed for each link.

Thus, as an RM cell traverses a path of a session, it encounters numerous MACR values, and the ER field maintains the lowest encountered MACR, or the value of 2 times CCR, whichever is smaller.

In implementations that are suited for TCP networks (wherein the active elements between links are called routers, rather than switches, and the information is contained in packets rather than cells), the MACR values are obtained through polling. TCP networks also have an Explicit Forward Congestion Indication (EFCI) bit. When a router encounters a cell with a CCR that is greater than the MACR computed for the just-traversed link, the bit is set.

When the RM reaches the end of the path, the value contained in the ER field is either 2 times CCR, or the smallest MACR computed in some link along the path that the RM cell traversed. Afek does not give a name to this value, so for sake of clarity, it is referred to herein as the Session Bottleneck Sensor (SBS). The RM cell is returned to the source, and the information this provides to the source is the SBS.

Thus, in non-TCP embodiments, the information that is available to the source is the SBS, but **the function that employs the SBS to arrive at a new ACR is not disclosed or in any manner taught by Afek.**

In TCP embodiments, the information that is available to the source is the polled MACR values, and the value of the EFCI bit. It is taught that when the source receives a set EFCI bit, it reduces the rate of the source. However, **the function that employs the various polled MACR values to arrive at a new ACR is also not disclosed or in any manner taught by Afek.**

All that is disclosed regarding the new rate (col. 6, lines 29-35) is that:

The rates of sessions that are above Δ are reduced towards Δ , and the rates of sessions that are below may be increased. This mechanism reaches a steady state only when the unused capacity Δ is equal to the maximum rate of any session that crosses the link and all the sessions that are constrained by this link are at this rate.

Afek does not identify the link that relates to the Δ mentioned above but, presumably, it is the Δ of the link whose MACR eventually reached the source.

Correspondences Asserted by Examiner

Before proceeding with a detailed analysis, it is useful to establish what the Examiner believes to be the correspondences between the teaching of Afek and the terms employed in claim 1.

- In connection with the first step of the claim, which specifies a “session congestion measure,” the Examiner points to the EFCI bit. The Examiner fails to *explicitly* state what corresponds to the “session congestion measure,” but the Examiner’s argument effectively asserts that EFCI bit corresponds to the session congestion measure; i.e.,

EFCI bit = session congestion measure.

Consequence: This correspondence also **limits** the Examiner’s assertions to TCP embodiments. Therefore, when perusing the Afek disclosure for how a new rate is computed, for example, the perusal is for how a new rate is computed for TCP embodiments.

- In connection with the second step of the claim, which specifies evaluation a “session incremental reward function that is related to rate of said incoming traffic,” the Examiner only points to the MACR. In the response to applicant’s argument, the Examiner cites col. 8, lines 7-14, and col. 6, lines 25-41. However, the col. 8 passage addresses the effect of Δ on a link’s MACR, and the col. 6

passage does not address MACR at all. In the rejection, the Examiner cites the same col. 8, lines 7-14, and col. 7, lines 63-65. The col. 7 citation merely provides the equation for calculating the MACR of a link. The Examiner fails to *explicitly* state what corresponds to the “session incremental reward function,” but based on the Examiner’s remarks, it appears that the Examiner equates the MACR of a link to the “session incremental reward function;” i.e.,

$$\text{MACR of a link} = \text{session incremental reward function} \\ \text{that is related to rate of said incoming traffic}$$

Consequence: In TCP embodiments, the source polls the switches and obtains the MACR values of a plurality of links. It appears that Afek uses the minimum MACR received during polling.

- In connection with the third step of the claim, which specifies a “global network cost function which combines cost functions assigned to said sessions and congestion cost functions assigned to said links,” the Examiner speaks of minimization of changes in MACR and changes in link utilization. Here, too, the Examiner fails to *explicitly* state what corresponds to the “global network cost function.” The Examiner also has not identified any function that is used as a guide in computing new rates. Applicant believes that the function that yields the SBS value is the function in Afek that controls the computation of new cell rates, but that is NOT what the Examiner is asserting, and even if the Examiner did make that assertion, it would not comport with the explicit limitation of the “global network cost function,” as demonstrated below. In TCP embodiments, as stated above, it is a pure guess as to what information (other than the EFCI bit) is used in order to compute the new rate – unless one speaks of the set of polled MACR values.
- Separately, the Examiner asserts that the MACR is the session rate (remarks of current Office action, at page 4, end of third paragraph. First, it is not clear whether the Examiner is referring to the MACR of links, or to the SBS. Secondly, the Examiner cannot claim that the MACR is the session rate, and also the “session incremental reward function.” Third, the MACR is not equal to the ACR.

Based on the above, applicant respectfully disagree that the correspondences set forth by the Examiner are valid, and also disagrees with the assertions made that new rates are calculated to minimize MACR and minimize changes in link utilization.

Discussion of claims addressed in “Response to Arguments”

Considering the first step of claim 1, as stated above, the correspondence asserted by the Examiner limits the Examiner’s assertions to TCP embodiments.

Considering the second step of claim 1, it is perfectly clear that the MACR, which is a function of Δ , the unused capacity of a **link**, is a parameter of a link (as stated above), and it depends on all of the sessions that share the link. Hence, the MACR of a link – which is the correspondence that the Examiner apparently asserts – cannot be a **session incremental reward function** – which is what claim 1 specified. The minimum MACR from the set of MACR values that the source obtains through polling also does not form a session reward function because that MACR value depends on all of the other sessions that share the link of that MACR and, therefore, it does not characterize the **session**. Additionally, an “incremental reward function” specified in the second clause of claim 1 is a function that reflects the **incremental** changes in a reward function. In other words, the session incremental reward function must be the derivative of another function, which is the session reward function. It is clear from the equations defining MACR that the MACR is not a derivative of another function¹

Considering the third step of claim 1, it specifies:

evaluating a new rate of said incoming traffic that moves said rate of said incoming traffic in a direction that minimizes a global network cost function which combines cost functions assigned to said sessions and congestion cost functions assigned to said links.

Clearly, in order for claim 1 to be anticipated, the reference must teach a global network cost function, and that function must be one that combines cost functions assigned to sessions and congestion cost function assigned to links. Also clearly, in order for claim 1 to be anticipated, the reference must teach that a new rate is evaluated is such that moves the rate in a direction that minimizes the identified cost function.

¹ In non-TCP embodiments, the MACR of a link is also is not a session reward function, for the same reason.

In rejecting claim 1, and in connection with this third step thereof, the Examiner states “[new flow rates are calculate to minimize the changes in MACR (cost functions assigned to sessions) and changes in link utilization (congestion cost functions assigned to links) (col. 8, line 25-col. 9, line 19)].”

1. This statement asserts that col. 8, lines 25 – col. 9, line 19 teaches that new flow rates are calculated to minimize changes in MACR. Respectfully, applicant sees no such teaching. None of the paragraphs address minimization of changes in MACR.
2. This statement also asserts that new flow rates are calculated to minimize changes in link utilizations. Respectfully, applicant sees no such teaching. None of the paragraphs address minimization of changes in link utilization.
3. This statement does NOT identify any global network cost function that is to be minimized. Indeed, the cited passage does not describe any function that calculates flow rates, and certainly it does not describe a function that is a “global network” function.

Analyzing the passage cited by the Examiner, it is noted that it contains five paragraphs. For sake of clarity, the analysis of the passage, which follows, handles whole paragraphs as units.

The first paragraph states:

If ACR exceeds MACR, it is possible that sessions for which the current rate is larger than MACR will cause the value of MACR to be decreased and hence to be underestimated. For example, assume a link in which two sessions are restricted and the bandwidth is 12. If one of those sessions is transmitting consistently at a rate of 8, a stable state can be achieved in which both MACR and the rate of the second session are set to 2. To avoid this phenomenon, which might cause unfairness, some preferred embodiments of the present invention, when computing Δ , treat the load caused by the sessions for which the indicated rate exceeds MACR as though it were exactly MACR. In other words, in the counting of the cells arriving at the output port of each link, a cell whose ACR exceeds MACR is counted as only MACR/ACR cells. For example, if MACR equals 4 and a session has a rate of 8, then only half the cells transmitted by that session are counted in the computation of Δ .

This paragraph discusses the notion of counting fewer than all of the cells that arrive at a node. There is no discussion of actually computing MACR, there is no discussion of

sessions, and there is no notion of a global network cost function. Moreover, other than the fact that MACR is related to link utilization, there is no discussion of link utilization, and certainly no discussion of minimizing changes in link utilization. In short this paragraph supports none of the Examiner's assertions.

The second paragraph states:

In the preferred embodiment of the present invention shown in FIG. 1, there are two weighting coefficients, α_{inc} and α_{dec} . These are provided to avoid sensitivity to queue length. If the same weighting coefficient is used regardless of the size of the queue and the rate of change of the queue, then, if many sessions pass through the same link, the session rates may suffer large oscillations and never converge, and the queue length may grow without bound. To avoid this, α_{inc} is used when Δ is greater than the prior MACR, and α_{dec} is used when Δ is less than or equal to the prior MACR. Moreover, the actual values of the weighting coefficients depend on the queue length. When the queue length is relatively small, α_{inc} is large and α_{dec} is small. This shortens the convergence time of sessions and decreases the period of time in which the link is underutilized. When the queue length is large, α_{dec} is large and α_{inc} is small, to decrease the queue length and prevent large delays and data loss. FIG. 2 contains an example of a table for computing α_{inc} and α_{dec} , based on a queue_threshold parameter and a base coefficient α .

This paragraph addresses the coefficients α_{inc} and α_{dec} , which (from other passages it is known) are used for computing values of MACR in a link. More particularly, this passage teaches when α_{inc} is used, and when α_{dec} is used. It also teaches that the values of α_{inc} and α_{dec} aim to reduce (note, not minimize) the link's queue length. What this passage fails to teach or suggest is how MACR is computed, or with what objective such a computation would be made. (the computation of MARC is presented elsewhere in the reference, but not in this paragraph.) Specifically, this paragraph does not teach or suggest minimizing changes to MACR (as asserted by the Examiner), and it does not teach or suggest minimizing changes in link utilization (as asserted by the Examiner). Indeed, this paragraph does not address sessions, and does not address links in the plural. Simply stated, this paragraph has no notion of a global network cost function that involves sessions (in plural) and links (in plural). This paragraph also does not support the assertion of minimizing changes to MACR, or minimizing changes in link utilizations.

The third paragraph states:

The first line of the pseudocode, in FIG. 1, in the block labeled "For every backward RM cell do:", implements the scheme described above for avoiding underutilization of the network in case only a few "heavy" sessions are using the network. The number compared to the value in the ER field of the backward traveling RM cell is not MACR itself, but MACR multiplied by a utilization factor f_u .

This paragraph addresses the MACR value that is sent back to a source (the SBS value) and in combination with the pseudocode of FIG. 1, it shows that the value that reaches a source is 2 times the current cell rate, or the MACR the link that has the lowest product of the MACR and a utilization_factor, whichever is lower. It teaches nothing about computing the actual new flows, or that the new flow computations minimize changes in MACR, or that the new flow calculations minimize changes in link utilization. Clearly there is no teaching of any global network cost function. Thus, this paragraph also does not support the Examiner's assertions.

The fourth paragraph states:

If the utilization factor f_u is significantly greater than 1, or if many "greedy" sessions are constrained on the link, then the value of MACR computed by the algorithm of FIG. 1 may be very oscillatory. The reason for this is that small changes in MACR are multiplied by f_u and subsequently affect all of the "greedy" sessions. FIG. 3 shows pseudocode for a method of stabilizing MACR, by computing its mean variation and modifying α_{inc} and α_{dec} accordingly. The mean variance of MACR is used in preference to the standard deviation of MACR because computing the mean variance does not require a square root computation.

This paragraph addresses some of the pseudocode shown in FIG. 1 and states that under certain conditions, when MACR changes happen to be small, the computation of MACR by using the FIG. 1 equations results in "very oscillatory" values of MACR. The paragraph teaches that, in such a case, the use of the equations described in FIG. 3 is preferable. What this paragraph fails to teach is the minimization of changes in MACR (as asserted by the Examiner) and, if fact, suggests that minimization of MACR is not necessarily a desired end result. This paragraph also does not teach minimization of changes in link utilizations. The passage certainly does not teach or suggest anything about a cost function that involves sessions (in plural) and links (in plural). Hence, this paragraph also does not support the Examiner's assertions.

The fifth paragraph states:

The usual approach to computing the mean variance of MACR is to do the following computations:

$$E := \text{MACR} - \alpha$$

$$D := D * (1 - h) + \text{ABS}(E) * h$$

where the weighting factor h is an inverse power of two, typically $1/16$. This approach, however, can not distinguish between the case where D has a large value due to an external change, such as an addition or a removal of a new session, and the case where the large variation stems from the fact that a small change in MACR causes a large change in link utilization. Only in the second case is it desirable to smooth the changes on MACR in order to achieve convergence.

This paragraph addresses the computation of the mean variance of MACR. It has nothing to do with a global network cost function that involves sessions (in plural) and links (in plural). Also, it does not teach computation of new flows, and it certainly does not teach computation of new flows that minimize changes to MACR or minimize changes to link utilizations.

In summary, it is respectfully submitted that the cited text between col. 8, line 18, and col. 8, line 19, does not support the Examiner's assertion that flow rates are calculated to minimize changes in MARC and changes in link utilization. Also, the cited text does not teach or suggest a “**global** network cost function which **combines** cost functions assigned to said sessions [in plural] and congestion cost functions assigned to said links [in plural]” (emphasis and parenthetical expressions supplied), as specified in claim 1. Further, it does not teach evaluating a new rate for incoming traffic that rate in the direction that minimizes this **global** cost function, as also specified in claim 1.

Lastly, it is noted that no global cost function is mentioned anywhere is Afek, in the sense that a cost function is global, it must be more than for just one session, and applicant's specification clearly gives the term “global” a global-over-the-sessions-in-the-network” meaning.

In light of the above analysis, it is respectfully submitted Afek does not anticipate that claim 1. Consequently, all of the remaining claims, which depend on claim 1, are also not anticipated by Afek.

In order to expedite prosecution, claim 1 is amended by explicitly limiting the session incremental reward function to be related to the rate of the session, and to traffic of no other session.

As for claim 5 (the next claim addressed in the Examiner's *Response to Arguments*), applicant made a number of arguments in favor of patentability. First, applicant asserted that Afek does not explicitly teach how a rate is computed. Therefore, applicant reasoned, Afek does not anticipate claim 5. The Examiner simply ignored this argument, and proceeded to challenge the next argument, which asserted that the claim specifies both an "incremental reward function" and a "session congestion measure" to compute a new rate, and that Afek does not use these two factors to compute a new rate. The Examiner asserts that Δ is used to compute the rates. First, that is not two factors. Second, respectfully, the reference does not say Δ is used to compute the rates. In fact, at col. 6, lines 57-64 Afek explicitly states that Δ is **not** used directly to compute rates, but that, rather, a weighted average of Δ is used to update the link's MACR. The Examiner points to col. 6, lines 24-35, but the cited passage only teaches that the rates are changed **toward Δ** .

What can be concluded from the above analysis of the teachings in Afek is that:

- (a) the calculation of a MACR of a link is related to Δ ,
- (b) the most heavily used link in the path of a session, which yields the minimum MACR of the links in the path of the session (the SBS), is sent back to the source in non-TCP embodiments,
- (c) in TCP embodiments MACR values of the path are obtained by the source through polling, and
- (d) a new flow that is calculated causes the new rate to be changed in a direction toward Δ .

Since the correspondence chosen by the Examiner restricts consideration to TCP embodiments (from the standpoint of anticipation), all that can be said is that the MACR values obtained through polling and the EFCI bit are the parameters that are available to the source for computing a new rate. Hence, the Examiner's assertion that Δ is used to compute the rates cannot be correct. Moreover, Δ is neither the incremental reward function (which the Examiner asserts to be the MACR of a link), nor is it the session

congestion measure (which the Examiner asserts to be the EFCI bit). Both the incremental reward function and the session congestion measure are needed to compute a rate in accordance with the subject claims. The Examiner has not shown use of the two factors as specified in claim 1. Nor has the Examiner shown that the computation works toward minimizing a cost function (only that it works to approach Δ).

Indeed, precisely how the minimum MACR value and the EFCI bit are used to compute the new rate is not known.

Applicants respectfully submit, therefore, that the Examiner erred in rejecting claim 5, because applicant's argument should have been deemed persuasive.

In connection with claim 6, the Examiner challenges applicant's interpretation of the term "receiving end of said session," asserting that "there is nothing in the claim that defines the receiving end other than a router." Applicant respectfully disagrees. A session is a term used in connection with a connection from a source to a destination. The connection traverses links and routers. Both the links and the routers are intermediate points of the session. There is only one transmitting end of a session -- that being the source that sends out packets--, and there is only one receiving end of a session -- that being the destination. Routers are intermediate points that can be said to have both receiving and transmitting elements. It is noted that applicant's sense of the term "receiving end" as it applies to claim 1 is perfectly consistent with the use of this term throughout the specification.

Further, the fact that the congestion field of applicant's specification is incremented in routers cannot be equated with the notion that rates are computed in the routers, and if anything, teaches away from such a conclusion. Moreover, and equally importantly, the Examiner's comments suggest that applicant's claim allows the interpretation that new rates are computed within routers, **and** that Afek also computes new rates in the routers. Respectfully, there is no support for an assertion that Afek computes new rates in the routers. The Examiner cites col. 6, lines 34-47, but the cited passage merely discusses how to evaluate Δ . It says absolutely nothing about computing rates in routers.

Applicant respectfully submits, therefore, that the Examiner erred in rejecting claim 6, because applicant's argument should have been deemed persuasive.

In connection with claim 8, the Examiner asserts that Δ “is used to constrain the rates of sessions by a multiple of Δ ,” citing col. 6, lines 57-67, and further asserts that this corresponds to the “additive factor” of claim 8. Applicant respectfully traverses on two grounds. First, the cited passage states that Δ is NOT used directly to constrain rates (contrary to the Examiner’s apparent assertion) and that a weighted average of Δ is effectively stored in MACR of the link. The passage then states that instead of using Δ , one might use Δ' , which is a multiple of Δ . That does not translate to a statement that rates of sessions are constrained by a multiple of Δ . Second, while the effect of Δ on MACR of a link may be said to be additive (not from the cited passage, but from the FIG. 1 equation

$$\text{MACR} := \text{MACR}(1 - \alpha_{\text{inc}}) + \Delta \alpha_{\text{dec}}$$

it is known that the MACR is NOT used to compute rates, so this additive relationship is not wholly relevant. It is true that the MACR of a link has the potential of affecting the rate computed by a source, but we know nothing of the relationship between the set of MACR values polled by a source and the rate computed by the source. Lastly, as for the notion of using a multiple of Δ instead of Δ , it is respectfully submitted that it would be contrary to logic to suggest that “multiple” suggests “additive.”

Applicant respectfully submits, therefore, that the Examiner erred in rejecting claim 8, because applicant’s argument should have been deemed persuasive.

Regarding claim 34, the Examiner challenges applicant’s argument by asserting that col. 9, lines 56-57 and/or col. 11 lines 5-11 disclose “that the value of the ER field in RM cells is equated to MACR.” Applicant respectfully submits that the cited passages disclose no such thing. The col. 9, lines 56-57 passage addresses a consequence where “in some unfortunate cases ...” (col. 9, line 47) “all rates are restricted to this link are allowed to get a rate equal to MACR, or to $\text{MACR} * f_u$.” It’s not clear which link’s MACR the reference discusses, but it is presumed that the references discusses the MACR of the link that is most heavily utilized. The fact that this is an unfortunate occurrence demonstrates that the rate was not purposely set to MACR; i.e., that Afek does not contemplate the equation

$$\text{Allowable Cell Rate} = \text{MACR}.$$

Rather, whatever function is used to compute the rate resulted in the rate being equal to MACR or $\text{MACR} \cdot f_u$ of that link. To give an example, the fact that the price of a luxury car in New York has risen to the price of a house in rural Canada does not mean that the price of houses in rural Canada is determined by the price of luxury cars in New York.

Applicant respectfully submits, therefore, that the Examiner erred in rejecting claim 34, because applicant's argument should have been deemed persuasive.

This ends the discussion of claims rejected under 35 USC 102 that were addressed in the Examiner's Response to Arguments. The following discusses the claims rejected under 35 USC 103.

Regarding claim 2, the Examiner challenges applicant's arguments by asserting, with respect to the Mitra reference, that "First, eq-n 15 a gradient of the subnetwork revenue with respect to capacity of the link which is indeed a cost function or a 'session incremental reward function'". Applicant has admitted that equation 15 is a revenue sensitivity equation with respect to link capacity. Applicant also noted that it is not a function of any traffic rate. To the extent that the Examiner asserts that equation 15 is a cost function, if it is to be equated to the session incremental reward function then it must be "the negative of a derivative, with respect to rate of said incoming traffic, of said ones of said costs functions assigned to said sessions." In other words, equation 15 should be of the form:

$$= - \frac{\partial[\text{a cost function}]}{\partial[\text{traffic rate}]}$$

However, equation 15 is

$$= - \sum \sum \frac{\partial L_s}{\partial C_l} v_{sl}; r \left(e_{sr} - \sum c_{sk} \right),$$

which is of a different form. When focusing on the derivative $-\frac{\partial L_s}{\partial C_l}$, it can be said to be

the negative of a derivative, with respect to link capacity, of function L_s , which is a link loss probability function (see col. 13, line 62). The link loss probability function is **not** a cost function that is a member of the cost functions of sessions that contribute to the global network cost function. Hence, equation 15 of Mitra does not correspond to the "session incremental reward function" of claim 2.

The Examiner also asserts that”

Secondly, Afek discloses that the flow control algorithm to regulate traffic and maximize bandwidth allocation (col. 7, lines 51-54) and Mitra discloses optimizing bandwidth by determining traffic rate to be offered and the allocations of bandwidth to respective links (abstract, col. 7, lines 62-66) which are analogous are.

Respectfully, this assertion fails to make sense. First, regarding Afek, the clause appears to be missing something, unless the phrase “Afek discloses that the flow” is replaced by “Afek discloses a flow.” In such a case, however, the statement fails to teach anything that is relevant to claim 2. Second, regarding Mitral, there are obviously some words missing from the end of the statement (since it ends with the word “are”). As for what Mitra does, in fact, teach, the Abstract (cited by the Examiner) clearly states that what is solved is the joint problem of optimal routing and bandwidth allocation. The optimization undertaken maximizes the revenue (see col. 17, line 61).

More importantly, the Examiner asserted in connection with claim 1 that the MACR of a link in Afek is the session incremental reward function. The MACR is carefully defined in Afek. In connection with claim 2, the Examiner notes the Mitra has a cost function with a negative derivative, and apparently suggests that the cost function of Mitra should replace the MACR of a link in Afek. Applicant respectfully submits that there is no motivation for doing so, and that the Afek system will simply not work as intended if the MACR of Afek is arbitrarily replaced with some other function that is constrained to be “the negative of a derivative,” generally, or “the negative of a derivative with respect to rate of said incoming traffic, of said one of said cost functions assigned to said sessions,” in particular.

In short, applicant respectfully submits that the Examiner erred in rejecting claim 2.

With respect to claim 3, the Examiner challenges applicant’s arguments by asserting that “the capacity cost of the link as shown in eq-n 3 are congestion cost functions of links.” Equation 3 states, effectively, that the sensitivity of subnetwork revenue is related to the negative of a sum of capacity costs (i.e., related to $-\sum c_{s,l}$).

Stated in other words, the derivative of the revenue, W, is related to the negative of a sum of capacity costs. Claim 3, in contradistinction, specifies that the session congestion

measure is a derivative with respect to rate of traffic of a sum of congestion cost functions; i.e.,

$$\text{session congestion measure} = \frac{\partial[\text{sum of cost functions}]}{\partial[\text{traffic rate}]}$$

Clearly, equation 3 of Mitra does not correspond to the congestion measure of claim 3.

More importantly, the Examiner asserted in connection with claim 1 that the EFCI bit in Afek is the session congestion measure. This bit is set in accordance with a specific procedure. In connection with claim 3, the Examiner notes equation 3 of Mitra, and apparently asserts that equation 3 should be substituted for the EFCI bit. Applicant respectfully submits that there is no motivation for doing so, and that the Afek system will simply not work as intended if the EFCI bit were replaced by the Equation 3 function (even if it did meet the limitations of claim 3, which it does not).

Thus, it is respectfully submitted that the Examiner erred in rejecting claim 3.

With respect to claim 4, the Examiner challenges applicant's arguments and asserts that "the link capacities may become so great (col. 6, lines 38-39) and a threshold is used to limit the upper bound of the link capacity (col. 12, lines 24-27)." This is the same statement that applicant struggled to understand in responding to the last Office action. The statement is unclear. Is the Examiner asserting that the passage in col. 6, lines 38-39 teaches that "the link capacities may become so great" and that the passage in col. 12, lines 24-27 teaches that "a threshold is used to limit the upper bound of the link capacity increment?" If so, applicant is at a loss. In the previous Office action applicant quoted the cited language in col. 6, and pointed out what it does teach. It does NOT teach anything about link capacities being small, large, or "so great." As for the col. 12 language, it speaks of a lower bound, an upper bound on link-capacity increments, and thresholds "that are used for testing the convergence of network revenue W." It does not disclose assigning a very large congestion cost function value (or cost) when a link load exceed what is considered to be the chosen maximum permissible link load.

It is respectfully submitted that neither the Examiner's comments nor the passages cited by the Examiner refute applicant's arguments, made in the previous office action and made now and, accordingly, it is respectfully submitted that the Examiner erred in rejecting claim 4.

With respect to claim 35, the Examiner points to the link loss probabilities mentioned in Mitra. It should be remembered that claim 35 simply specified that the session congestion measure is based on probability of packet loss experienced at the receiving end. While rejecting claim 1, the Examiner asserted a correspondence between the EFCI bit and the session congestion measure. That is a single bit. In connection with claim 3, the Examiner suggested that equation 3 of Mitra should be substituted for the EFCI bit. As indicated above, there is no motivation for such a substitution. In connection with claim 35, it is not clear what the Examiner asserts to correspond to the session congestion measure, but the remarks generally mention link-loss probability estimates. Respectfully, just because Mitra's invention contains considerations of link loss probabilities does not suggest that the EFCI bit of Afek should be replaced with a measure that is "based on probability of packet loss experienced at said receiving end." There is simply no connection between them. Hence, it is respectfully submitted that the Examiner erred in rejecting claim 35.

With respect to claim 12, the Examiner challenges applicant's statement that FIG. 3 of the Szentesi reference "relates to revenue versus traffic for the **entire network**. It is not an incremental reward function for a **session**" (emphasis in original), asserting that "FIG. 3 represents the relation of traffic flow and revenue (i.e., session incremental reward function)." Applicant respectfully traverses. First, FIG. 3 is described by Szentesi as a figure that "shows the operating principle underlying the operation of the subject algorithm using the exemplary **network** throughput curve of FIG. 2 as a template" (emphasis supplied). Hence, it is respectfully submitted that applicant's assertion in the previous Office response was correct. Second, the correctness of applicant's assertion is not diminished by the correctness of the Examiner's assertion that "FIG. 3 represents the relation of traffic flow and revenue." Both statements are correct. Third, the fact that FIG. 3 represents the relation of traffic flow and revenue, does NOT lead to the conclusion that FIG. 3 describes a session incremental function. It is neither an incremental function, nor is it related to a session. One might say that a *derivative* of the FIG. 3 function would describe an incremental function. However, a derivative of the FIG. 3 curve still would not make it a session incremental function. Lastly, even if it were a session incremental function, it is noted that the derivative of the FIG. 3 function

is not a “positive, decreasing” function with respect to session , because beyond load t_m , the derivative is clearly negative. Even the FIG. 3 function itself does not meet the “positive, decreasing” limitation of claim 12 (since it’s increasing at loads smaller than t_m). In short, it is respectfully submitted that the Examiner erred in rejecting claim 12.

As for claim 14, it addresses link cost functions (in contrast to the session incremental reward function of claim 12). Here, too, the Examiner points to FIG. 3 of the Szentesi reference. Respectfully, it’s unclear how the Examiner can assert that FIG. 3 describes the session incremental reward function, and also assert that FIG. 3 describes the link cost function. It seems that the Examiner should be bound to one assertion, and required to stay with that assertion. Applicant believed that FIG. 3 describes neither, as explained above. Moreover, claim 14 specifies a “positive, increasing” function. FIG. 3 is not a positive, increasing, function because it decreases beyond load t_m . Applicant respectfully submits, therefore, that the Examiner erred in rejecting claim 14.

Regarding claim 15, it addresses the derivative of link cost functions. Again the Examiner points to FIG. 3. Applicant respectfully adopts the argument employed in connection with claim 14.


Discussion of claims NOT addressed in “Response to Arguments”

It is noted that claims that were rejected in this Office action and not addressed above were rejected in the previous Office action with identical reasons. Applicant respectfully adopts the arguments made in the previous Office response, noting that the Examiner offered no “response to arguments” in connection with these claims. In other words, it appears that the Examiner had no rebuttal to applicant’s arguments pointing to the patentability of these claims.

In light of the above amendments and remarks, applicant respectfully submits that all of the Examiner's rejections and objections have been overcome. Reconsideration and allowance are respectfully solicited.

Dated: 7/20/04

Respectfully,
Jamaloddin S Golestani

By 
Henry T. Brendzel
Reg. No: 26,844
Phone (973) 467-2025
Fax (973) 467-6589
email brendzel@comcast.net